



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
-----------------	-------------	----------------------	---------------------	------------------

10/804,976

03/18/2004

Amit Ganesh

50277-2502

7843

7590

10/11/2006

McDERMOTT WILL & EMERY
600 13th Street, N.W.
Washington, DC 20005-3096

EXAMINER

FLEURANTIN, JEAN B

ART UNIT

PAPER NUMBER

2162

DATE MAILED: 10/11/2006

Please find below and/or attached an Office communication concerning this application or proceeding.

Office Action Summary

Application No.

10/804,976

Applicant(s)

GANESH ET AL.

Examiner

JEAN B. FLEURANTIN

Art Unit

2162

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

- 1) ☒ Responsive to communication(s) filed on 18 March 2004.
- 2a) ☐ This action is **FINAL**. 2b) ☒ This action is non-final.
- 3) ☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

- 4) ☒ Claim(s) 1-40 is/are pending in the application.
- 4a) Of the above claim(s) _____ is/are withdrawn from consideration.
- 5) ☐ Claim(s) _____ is/are allowed.
- 6) ☒ Claim(s) 1-40 is/are rejected.
- 7) ☐ Claim(s) _____ is/are objected to.
- 8) ☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

- 9) ☐ The specification is objected to by the Examiner.
- 10) ☐ The drawing(s) filed on _____ is/are: a) ☐ accepted or b) ☐ objected to by the Examiner.
Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).
- 11) ☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

- 12) ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
- a) ☐ All b) ☐ Some * c) ☐ None of:
1. ☐ Certified copies of the priority documents have been received.
 2. ☐ Certified copies of the priority documents have been received in Application No. _____.
 3. ☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).
- * See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

- | | |
|---|---|
| 1) <input checked="" type="checkbox"/> Notice of References Cited (PTO-892) | 4) <input type="checkbox"/> Interview Summary (PTO-413)
Paper No(s)/Mail Date. _____ |
| 2) <input type="checkbox"/> Notice of Draftsperson's Patent Drawing Review (PTO-948) | 5) <input type="checkbox"/> Notice of Informal Patent Application |
| 3) <input checked="" type="checkbox"/> Information Disclosure Statement(s) (PTO/SB/08)
Paper No(s)/Mail Date <u>7/19/4</u> . | 6) <input type="checkbox"/> Other: _____ |

DETAILED ACTION

1. This is in response to the application filed on 3/18/04.

Claims 1- 40 are presented for examination.

Drawings

The Drawings submitted on 3/18/04 are acknowledged.

The document submitted on 06/21/04 has been entered and considered.

Information Disclosure Statement

The information disclosure statement (IDS) submitted on 07/19/04. The submission is in compliance with the provisions of 37 CFR 1.97. Accordingly, the information disclosure statement is being considered by the examiner.

Claim Rejections - 35 USC § 101

35 U.S.C. 101 reads as follows:

Whoever invents or discovers any new and useful process, machine, manufacture, or composition of matter, or any new and useful improvement thereof, may obtain a patent therefor, subject to the conditions and requirements of this title.

Claims 1-20 are rejected under 35 U.S.C. 101 because the claimed invention is directed to non-statutory subject matter.

As set forth in MPEP 2106:

Products may be either machines, manufactures, or compositions of matter.

A *machine* is "a concrete thing, consisting of parts or of certain devices and combinations of devices." *Burr v. Duryee*, 68 U.S. (1 Wall.) 531, 570 (1863).

As per claim 1,

Claim 1, in view of the above cited MPEP section, "a method for performing database recovery after a crash of an instance of a database, wherein multiple transactions were active when the instance crashed, the method comprising the steps of identifying a plurality of dead transactions; determining statistical data about said plurality of dead transactions; determining that a particular number of recovery servers should be used to recover said plurality of dead transactions based on the statistical data; and recovering said plurality of dead transactions using said particular number of recovery servers," does not produce any useful and tangible result.

As per claim 16,

Claim 16, in view of the above cited MPEP section, "a computer-readable medium carrying one or more sequences of one or more instructions for performing database recovery after a crash of an instance of a database, wherein multiple transactions were active when the instance crashed, wherein the execution of the one or more sequences of one or more instructions by one or more processors causes the one or more processors to perform the steps of identifying a plurality of dead transactions; determining statistical data about said plurality of dead transactions; determining that a particular number of recovery servers should be used to recover said plurality of dead transactions based on the statistical data; and recovering said plurality of dead transactions using said particular number of recovery servers," does not produce any useful and tangible result.

As per claim 31,

Claim 31, in view of the above cited MPEP section, "a system for performing database recovery after a crash of an instance of a database, wherein multiple transactions were active when the instance crashed, the system comprising a memory; one or more processors coupled to the memory; and a set of computer instructions contained in the memory, the set of computer instructions including computer instructions which when executed by one or more processors, cause the one or more processors to perform the steps of identifying a plurality of dead transactions; determining statistical data about said plurality of dead transactions; determining that a particular number of recovery servers should be used to recover said plurality of dead transactions based on the statistical data; and recovering said plurality of dead transactions using said particular number of recovery servers," does not produce any useful and tangible result.

And, all dependent claims are rejected under the analysis of claims 1 and 13, and these claims are rejected on that basis.

Claim Rejections - 35 USC § 103

The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

Claims 1-40 are rejected under 35 U.S.C. 103(a) as being unpatentable over "Aries-RRH: Restricted Repeating of History in THE Aries Transaction Recovery Method – pages 718-727, 1991" issued to Mohan et al., ("Mohan"), submitted by the applicants in view of US Pat. No. 5,544,359 issued to Tada et al., ("Tada").

As per claim 1, Mohan discloses "a method for performing database recovery after a crash of an instance of a database" (i.e., missing updates transaction are redone; see page 720, col. 1, lines 1-8), "wherein multiple transactions were active when the instance crashed" (i.e., repeating failures during restart recovery or rollbacks; see page 719, col. 1, last paragraph lines 10-14), the method comprising the steps of:

"identifying a plurality of dead transactions" (i.e., identifying missing updates, loser transactions; see page 720, col. 1, paragraph 1, lines 4-11); and

"recovering said plurality of dead transactions using said particular number of recovery servers" (i.e., tracking of the progress of rollback; see page 719, col. 2, first paragraph, lines 10-17).

Mohan fails to explicitly disclose determining statistical data (statistic information) about said plurality of dead transactions; determining that a particular number of recovery servers should be used to recover said plurality of dead transactions based on the statistical data (statistic information). However, Tada discloses a similar method, recovery operation on the database (see Tada col. 13, lines 16-22) determining statistical data (statistic information) about said plurality of dead transactions; determining that a particular number of recovery servers should be used to recover said plurality of dead transactions

based on the statistical data (statistic information) (see Tada col. 13, line 65 to col. 14, line 2), further, Tada discloses the amount of occupation of the hlf buffer on nonvolatile memory. It would have been obvious to a person of ordinary skill in the art at the time the invention was made to modify the method of Mohan by determining statistical data (statistic information) about said plurality of dead transactions; determining that a particular number of recovery servers should be used to recover said plurality of dead transactions based on the statistical data (statistic information) as disclosed by Tada (see Tada col. 13, lines 31-45). Such a modification would allow the method of Mohan to provide managing the input data for the database recovery independently of each other database by means of acquiring the log data by every database, in addition, the log data of the database to be recovered can be retrieved for a shorter time, allowing reduction of the time for the database recovery (see Tada col. 18, lines 35-45), therefore, improving the transaction performance and the accuracy of the parallel transaction recovery.

As per claim 2, Mohan discloses "recovering said plurality of dead transactions is performed by executing the particular number of recovery servers in parallel" (i.e., performing in parallel processing the new transactions; see page 722, col. 1, first paragraph, lines 7-9 and page 720, col. 1, second paragraph, lines 8-12).

As per claim 3, Mohan discloses "identifying the plurality of dead transactions includes the step of maintaining a working list" (i.e., updating and a log record is written; see page 719, col. 1, third paragraph, lines 1-5).

Mohan fails to explicitly disclose determining statistical data (statistic information) includes the step of determining statistical data (statistic information) based on the list dead transactions. However, Tada discloses a similar method, recovery operation on the database (see Tada col. 13, lines 16-22) determining statistical data (statistic information) includes the step of determining statistical data (statistic information) based on the list dead transactions (see Tada col. 13, line 65 to col. 14, line 2), further, Tada discloses the amount of occupation of the hlf buffer on nonvolatile memory. It would have been obvious to a person of ordinary skill in the art at the time the invention was made to modify the method of Mohan by

determining statistical data (statistic information) includes the step of determining statistical data (statistic information) based on the list dead transactions as disclosed by Tada (see Tada col. 13, lines 31-45). Such a modification would allow the method of Mohan to provide managing the input data for the database recovery independently of each other database by means of acquiring the log data by every database, in addition, the log data of the database to be recovered can be retrieved for a shorter time, allowing reduction of the time for the database recovery (see Tada col. 18, lines 35-45), therefore, improving the transaction performance and the accuracy of the parallel transaction recovery.

As per claim 4, in addition to claim 3, Mohan further discloses "locating a rollback segment, wherein the rollback segment contains a transaction table that contains entries associated with dead transactions" (i.e., missing updates transactions are redone during the redo pass; see page 720, col. 2, second paragraph, lines 1-5 and Fig. 3);

"scanning the transaction table to identify the dead transactions; and storing the identity of the dead transactions in the working list" (in light the specification at paragraph [0117], the purposed of scanning the transaction table is for identifying entries being associating with dead transactions. The method for using a log sequence number on every database page to track the page's state, in which every page update and log record is written is disclosed by Mohan page 719, third paragraph).

As per claim 5, in addition to claim 1, Mohan further discloses "maintaining a block count, wherein the block count identifies the number of undo blocks that are associated with a particular transaction" (i.e., tracking of the progress of rollback; see page 719, col. 2, first paragraph, lines 10-17); and the step of "wherein the total number of undo blocks is based on the block count associated with the dead transactions identified in the working list" (i.e., updating and a log record (block) is written; see page 719, col. 1, third paragraph, lines 1-5).

As per claim 6, Mohan substantially discloses the claimed subject matter, except determining statistical data (statistic information) includes the step of determining statistical data (statistic information)

based on the number of dead transactions that are identified in the working list. However, Tada discloses a similar method, recovery operation on the database (see Tada col. 13, lines 16-22) determining statistical data (statistic information) includes the step of determining statistical data (statistic information) based on the list dead transactions (see Tada col. 13, line 65 to col. 14, line 2), further, Tada discloses the amount of occupation of the hlf buffer on nonvolatile memory. It would have been obvious to a person of ordinary skill in the art at the time the invention was made to modify the method of Mohan by determining statistical data (statistic information) includes the step of determining statistical data (statistic information) based on the list dead transactions as disclosed by Tada (see Tada col. 13, lines 31-45). Such a modification would allow the method of Mohan to provide managing the input data for the database recovery independently of each other database by means of acquiring the log data by every database, in addition, the log data of the database to be recovered can be retrieved for a shorter time, allowing reduction of the time for the database recovery (see Tada col. 18, lines 35-45), therefore, improving the transaction performance and the accuracy of the parallel transaction recovery.

As per claim 7, Mohan discloses "determining that the particular number of recovery servers should be used includes the step of determining that the particular number of recovery servers should be used based on a max_parallelism threshold value, wherein the max_parallelism threshold value provides an upper limit for the number of recovery servers to be used" (i.e, continually taking the maximum of the LSN (maximum parallel threshold) of the next log record to be processed for each of the yet to be completely undone loser transaction until no loser transaction remains to be undone; see page 720, col. 1, second paragraph, lines 1-7 and Fig. 1).

As per claim 8, Mohan discloses "determining the max_parallelism threshold value based on a user input value" (i.e, continually taking the maximum of the LSN (maximum parallel threshold) of the next log record to be processed for each of the yet to be completely undone loser transaction; see page 720, col. 1, second paragraph, lines 1-7 and Fig. 1).

As per claim 9, Mohan discloses "identifying a rollback segment that was previously owned by the crashed instance at the time of its crash" (i.e., updates rolling back; see page 720, col. 1, second paragraph, lines 1-2); and

"the crashed instance reacquiring ownership of the rollback segment after the crashed instance is restarted" (i.e., updating logged on stable storage but whose effects on the data base pages did not get reflected on disk before the system failure; see page 720, col. 1, first paragraph, lines 1-9).

As per claim 10, Mohan discloses "reacquiring ownership of the rollback segment includes the steps of identifying an instance that currently owns the rollback segment that was previously owned by the crashed instance at the time of its crash" (i.e., identifying the transactions that are active, and addressing most recently written log; see page 719, col. 2, last paragraph, lines 1-6);

"requesting the instance to release ownership of the rollback segment" (i.e., determining the list of transactions rolling back in the undo pass; see page 719, col. 2, last paragraph, lines 10-14); and

"the instance releasing ownership of the rollback segment in response to the request" (i.e., updates rolling back; see page 720, col. 1, second paragraph, lines 1-2).

As per claim 11, Mohan further discloses "identifying a rollback segment that is unowned, wherein the unowned rollback segment is not currently associated with any instance of the database" (i.e., identifying the transactions that are active, and addressing most recently written log; see page 719, col. 2, last paragraph, lines 1-6); and

"associating the unowned rollback segment with the crashed instance" (i.e., repeating failures during restarting recovery, which undo the same non record multiple times; see page 719, col. 2, last paragraph, lines 10-19), "wherein associating the unowned rollback segment with the crashed node causes the rollback segment to be owned by the crashed instance" (i.e., updates rolling back; see page 720, col. 1, second paragraph, lines 1-2).

As per claim 12, Mohan further discloses "recovering the plurality of dead transactions comprises the steps of: maintaining a working list, wherein the working list identifies a list of dead transactions for which recovery will be attempted" (i.e., updating and a log record is written; see page 719, col. 1, third paragraph, lines 1-5);

"selecting a dead transaction from the working list" (i.e., selecting missing updates, loser transactions; see page 720, col. 1, paragraph 1, lines 4-11);

"acquiring a rollback segment lock on a rollback segment, wherein the rollback segment is associated with a transaction table that contains an entry that corresponds to the dead transaction" (i.e., performing update of the log record when locking is in effect; see page 722, col. 1, last paragraph to col. 2, up to line 9 and Fig. 6);

"acquiring a transaction lock on a chain of undo, wherein the chain of undo contain change information associated with the dead transaction" (i.e., performing update of the log record when locking is in effect; see page 722, col. 1, last paragraph to col. 2, up to line 9 and Fig. 6);

"determining whether the dead transaction still needs to be recovered; and if the dead transaction still needs to be recovered, assigning the dead transaction to a recovery server" (i.e., tracking of the progress of rollback; see page 719, col. 2, first paragraph, lines 10-17).

As per claim 13, Mohan discloses "acquiring the transaction lock includes the step of a coordinator acquiring the transaction lock" (in light the specification at paragraph [0121], the purposed of acquiring the transaction lock is for identifying dead transactions that need to be retrieved. The method for performing update of the log record when locking is disclosed by Mohan see page 722, col. 2, paragraph (1), lines 1-3).

As per claim 14, Mohan further discloses "upon completing the recovery of the dead transaction, the recovery server signaling the coordinator to indicate it has completed the recovery of the dead transaction" (In light the specification at paragraph [0139], the purposed of signaling the coordinator is for

indicating recovery of dead transactions had completed. The method of rolling back, until no loser transaction remains to be undo is disclosed by Mohan page 720, col. 1, second paragraph, lines 1-7); and

“upon receiving the signal from the recovery server, the coordinator releasing its lock on the transaction” (In light the specification at paragraph [0127], page 36, lines 5-7, the purposed of releasing lock on the transaction is for performing recovery of dead transactions. The method of recovering, tracking changing is disclosed by Mohan page 719, col. 1, last paragraph, lines 1-3).

As per claim 15, in addition to claim 1, Mohan further discloses “assigning two or more dead transactions to a recovery server” (i.e., tracking of the progress of rollback; see page 719, col. 2, first paragraph, lines 10-17);

“associating a time slice value with the recovery server, wherein the time slice value is used by the recovery server to promote fairness during recovery of the two or more dead transactions; and recovering the two or more dead transactions using the time slice value” (In light the specification at paragraph [0151], the purposed of associating a time slice value with the recovery server is for recovering two or more dead transactions. The method for processing transaction continually until no loser (dead) transaction remains to be undone is disclosed by Mohan see page 720, col. 1, second paragraph, lines 1-10).

As per claim 16, Mohan discloses “a computer-readable medium carrying one or more sequences of one or more instructions for performing database recovery after a crash of an instance of a database” (i.e., missing updates transaction are redone; see page 720, col. 1, lines 1-8), “wherein multiple transactions were active when the instance crashed” (i.e., repeating failures during restart recovery or rollbacks; see page 719, col. 1, last paragraph lines 10-14), “wherein the execution of the one or more sequences one or more instructions by one or more processors causes the one or more processor to perform” (i.e., performing parallel processing new transactions; see page 722, col. 1, first paragraph, lines 7-9) the steps of:

"identifying a plurality of dead transactions" (i.e., identifying missing updates, loser transactions; see page 720, col. 1, paragraph 1, lines 4-11); and

"recovering said plurality of dead transactions using said particular number of recovery servers" (i.e., tracking of the progress of rollback; see page 719, col. 2, first paragraph, lines 10-17).

Mohan fails to explicitly disclose determining statistical data (statistic information) about said plurality of dead transactions; determining that a particular number of recovery servers should be used to recover said plurality of dead transactions based on the statistical data (statistic information). However, Tada discloses a similar method, recovery operation on the database (see Tada col. 13, lines 16-22) determining statistical data (statistic information) about said plurality of dead transactions; determining that a particular number of recovery servers should be used to recover said plurality of dead transactions based on the statistical data (statistic information) (see Tada col. 13, line 65 to col. 14, line 2), further, Tada discloses the amount of occupation of the hlf buffer on nonvolatile memory. It would have been obvious to a person of ordinary skill in the art at the time the invention was made to modify the method of Tada by determining statistical data (statistic information) about said plurality of dead transactions; determining that a particular number of recovery servers should be used to recover said plurality of dead transactions based on the statistical data (statistic information) as disclosed by Tada (see Tada col. 13, lines 31-45). Such a modification would allow the method of Mohan to provide managing the input data for the database recovery independently of each other database by means of acquiring the log data by every database, in addition, the log data of the database to be recovered can be retrieved for a shorter time, allowing reduction of the time for the database recovery (see Tada col. 18, lines 35-45), therefore, improving the transaction performance and the accuracy of the parallel transaction recovery.

As per claim 17, Mohan discloses "recovering said plurality of dead transactions is performed by executing the particular number of recovery servers in parallel" (i.e., performing in parallel processing the new transactions; see page 722, col. 1, first paragraph, lines 7-9 and page 720, col. 1, second paragraph, lines 8-12).

As per claim 18, Mohan discloses "identifying the plurality of dead transactions includes the step of maintaining a working list" (i.e., updating and a log record is written; see page 719, col. 1, third paragraph, lines 1-5).

Mohan fails to explicitly disclose determining statistical data (statistic information) includes the step of determining statistical data (statistic information) based on the list dead transactions. However, Tada discloses a similar method, recovery operation on the database (see Tada col. 13, lines 16-22) determining statistical data (statistic information) includes the step of determining statistical data (statistic information) based on the list dead transactions (see Tada col. 13, line 65 to col. 14, line 2), further, Tada discloses the amount of occupation of the hlf buffer on nonvolatile memory. It would have been obvious to a person of ordinary skill in the art at the time the invention was made to modify the method of Mohan by determining statistical data (statistic information) includes the step of determining statistical data (statistic information) based on the list dead transactions as disclosed by Tada (see Tada col. 13, lines 31-45). Such a modification would allow the method of Mohan to provide managing the input data for the database recovery independently of each other database by means of acquiring the log data by every database, in addition, the log data of the database to be recovered can be retrieved for a shorter time, allowing reduction of the time for the database recovery (see Tada col. 18, lines 35-45), therefore, improving the transaction performance and the accuracy of the parallel transaction recovery.

As per claim 19, in addition to claim 3, Mohan further discloses "locating a rollback segment, wherein the rollback segment contains a transaction table that contains entries associated with dead transactions" (i.e., missing updates transactions are redone during the redo pass; see page 720, col. 2, second paragraph, lines 1-5 and Fig. 3);

"scanning the transaction table to identify the dead transactions; and storing the identity of the dead transactions in the working list" (in light the specification at paragraph [0117], the purposed of scanning the transaction table is for identifying entries being associating with dead transactions. The method for using a log sequence number on every database page to track the page's state, in which every page update and log record is written is disclosed by Mohan page 719, third paragraph).

As per claim 20, in addition to claim 1, Mohan further discloses "maintaining a block count, wherein the block count identifies the number of undo blocks that are associated with a particular transaction" (i.e., tracking of the progress of rollback; see page 719, col. 2, first paragraph, lines 10-17); and the step of "wherein the total number of undo blocks is based on the block count associated with the dead transactions identified in the working list" (i.e., updating and a log record (block) is written; see page 719, col. 1, third paragraph, lines 1-5).

As per claim 21, Mohan substantially discloses the claimed subject matter, except determining statistical data (statistic information) includes the step of determining statistical data (statistic information) based on the number of dead transactions that are identified in the working list. However, Tada discloses a similar method, recovery operation on the database (see Tada col. 13, lines 16-22) determining statistical data (statistic information) includes the step of determining statistical data (statistic information) based on the list dead transactions (see Tada col. 13, line 65 to col. 14, line 2), further, Tada discloses the amount of occupation of the hlf buffer on nonvolatile memory. It would have been obvious to a person of ordinary skill in the art at the time the invention was made to modify the method of Mohan by determining statistical data (statistic information) includes the step of determining statistical data (statistic information) based on the list dead transactions as disclosed by Tada (see Tada col. 13, lines 31-45). Such a modification would allow the method of Mohan to provide managing the input data for the database recovery independently of each other database by means of acquiring the log data by every database, in addition, the log data of the database to be recovered can be retrieved for a shorter time, allowing reduction of the time for the database recovery (see Tada col. 18, lines 35-45), therefore, improving the transaction performance and the accuracy of the parallel transaction recovery.

As per claim 22, Mohan discloses "determining that the particular number of recovery servers should be used includes the step of determining that the particular number of recovery servers should be used based on a max_parallelism threshold value, wherein the max_parallelism threshold value provides

an upper limit for the number of recovery servers to be used" (i.e, continually taking the maximum of the LSN (maximum parallel threshold) of the next log record to be processed for each of the yet to be completely undone loser transaction until no loser transaction remains to be undone; see page 720, col. 1, second paragraph, lines 1-7 and Fig. 1).

As per claim 23, Mohan discloses "determining the max_parallelism threshold value based on a user input value" (i.e, continually taking the maximum of the LSN (maximum parallel threshold) of the next log record to be processed for each of the yet to be completely undone loser transaction; see page 720, col. 1, second paragraph, lines 1-7 and Fig. 1).

As per claim 24, Mohan discloses "identifying a rollback segment that was previously owned by the crashed instance at the time of its crash" (i.e., updates rolling back; see page 720, col. 1, second paragraph, lines 1-2); and

"the crashed instance reacquiring ownership of the rollback segment after the crashed instance is restarted" (i.e., updating logged on stable storage but whose effects on the data base pages did not get reflected on disk before the system failure; see page 720, col. 1, first paragraph, lines 1-9).

As per claim 25, Mohan discloses "reacquiring ownership of the rollback segment includes the steps of identifying an instance that currently owns the rollback segment that was previously owned by the crashed instance at the time of its crash" (i.e., identifying the transactions that are active, and addressing most recently written log; see page 719, col. 2, last paragraph, lines 1-6);

"requesting the instance to release ownership of the rollback segment" (i.e., determining the list of transactions rolling back in the undo pass; see page 719, col. 2, last paragraph, lines 10-14); and

"the instance releasing ownership of the rollback segment in response to the request" (i.e., updates rolling back; see page 720, col. 1, second paragraph, lines 1-2).

As per claim 26, Mohan further discloses "identifying a rollback segment that is unowned, wherein the unowned rollback segment is not currently associated with any instance of the database" (i.e., identifying the transactions that are active, and addressing most recently written log; see page 719, col. 2, last paragraph, lines 1-6); and

"associating the unowned rollback segment with the crashed instance" (i.e., repeating failures during restarting recovery, which undo the same non record multiple times; see page 719, col. 2, last paragraph, lines 10-19), "wherein associating the unowned rollback segment with the crashed node causes the rollback segment to be owned by the crashed instance" (i.e., updates rolling back; see page 720, col. 1, second paragraph, lines 1-2).

As per claim 27, Mohan further discloses "recovering the plurality of dead transactions comprises the steps of: maintaining a working list, wherein the working list identifies a list of dead transactions for which recovery will be attempted" (i.e., updating and a log record is written; see page 719, col. 1, third paragraph, lines 1-5);

"selecting a dead transaction from the working list" (i.e., selecting missing updates, loser transactions; see page 720, col. 1, paragraph 1, lines 4-11);

"acquiring a rollback segment lock on a rollback segment, wherein the rollback segment is associated with a transaction table that contains an entry that corresponds to the dead transaction" (i.e., performing update of the log record when locking is in effect; see page 722, col. 1, last paragraph to col. 2, up to line 9 and Fig. 6);

"acquiring a transaction lock on a chain of undo, wherein the chain of undo contain change information associated with the dead transaction" (i.e., performing update of the log record when locking is in effect; see page 722, col. 1, last paragraph to col. 2, up to line 9 and Fig. 6);

"determining whether the dead transaction still needs to be recovered; and if the dead transaction still needs to be recovered, assigning the dead transaction to a recovery server" (i.e., tracking of the progress of rollback; see page 719, col. 2, first paragraph, lines 10-17).

As per claim 28, Mohan discloses "acquiring the transaction lock includes the step of a coordinator acquiring the transaction lock" (in light the specification at paragraph [0121], the purposed of acquiring the transaction lock is for identifying dead transactions that need to be retrieved. The method for performing update of the log record when locking is disclosed by Mohan see page 722, col. 2, paragraph (1), lines 1-3).

As per claim 29, Mohan further discloses "upon completing the recovery of the dead transaction, the recovery server signaling the coordinator to indicate it has completed the recovery of the dead transaction" (In light the specification at paragraph [0139], the purposed of signaling the coordinator is for indicating recovery of dead transactions had completed. The method of rolling back, until no loser transaction remains to be undo is disclosed by Mohan page 720, col. 1, second paragraph, lines 1-7); and

"upon receiving the signal from the recovery server, the coordinator releasing its lock on the transaction" (In light the specification at paragraph [0127], page 36, lines 5-7, the purposed of releasing lock on the transaction is for performing recovery of dead transactions. The method of recovering, tracking changing is disclosed by Mohan page 719, col. 1, last paragraph, lines 1-3).

As per claim 30, in addition to claim 1, Mohan further discloses "assigning two or more dead transactions to a recovery server" (i.e., tracking of the progress of rollback; see page 719, col. 2, first paragraph, lines 10-17);

"associating a time slice value with the recovery server, wherein the time slice value is used by the recovery server to promote fairness during recovery of the two or more dead transactions; and recovering the two or more dead transactions using the time slice value" (In light the specification at paragraph [0151], the purposed of associating a time slice value with the recovery server is for recovering two or more dead transactions. The method for processing transaction continually until no loser (dead) transaction remains to be undone is disclosed by Mohan see page 720, col. 1, second paragraph, lines 1-10).

As per claim 31, Mohan discloses "a system for performing database recovery after a crash of an instance of a database" (i.e., missing updates transaction are redone; see page 720, col. 1, lines 1-8), "wherein multiple transactions were active when the instance crashed" (i.e., repeating failures during restart recovery or rollbacks; see page 719, col. 1, last paragraph lines 10-14), the system comprising:

"a memory" (i.e., DB2; see page 720, paragraph (3), lines 1-2 and Fig. 3);

"one or more processors coupled to the memory" (i.e., processing of transactions in the database; see page 719, col. 1, last paragraph, lines 1-5); and

"a set of computer instructions contained in the memory, the set of computer instructions including computer instructions which when executed by one or more processors" (i.e., performing parallel processing new transactions; see page 722, col. 1, first paragraph, lines 7-9), causes the one or more processor to perform the steps of:

"identifying a plurality of dead transactions" (i.e., identifying missing updates, loser transactions; see page 720, col. 1, paragraph 1, lines 4-11); and

"recovering said plurality of dead transactions using said particular number of recovery servers" (i.e., tracking of the progress of rollback; see page 719, col. 2, first paragraph, lines 10-17).

Mohan fails to explicitly disclose determining statistical data (statistic information) about said plurality of dead transactions; determining that a particular number of recovery servers should be used to recover said plurality of dead transactions based on the statistical data (statistic information). However, Tada discloses a similar method, recovery operation on the database (see Tada col. 13, lines 16-22) determining statistical data (statistic information) about said plurality of dead transactions; determining that a particular number of recovery servers should be used to recover said plurality of dead transactions based on the statistical data (statistic information) (see Tada col. 13, line 65 to col. 14, line 2), further, Tada discloses the amount of occupation of the hlf buffer on nonvolatile memory. It would have been obvious to a person of ordinary skill in the art at the time the invention was made to modify the method of Mohan by determining statistical data (statistic information) about said plurality of dead transactions; determining that a particular number of recovery servers should be used to recover said plurality of dead transactions based on the statistical data (statistic information) as disclosed by Tada (see Tada col. 13,

lines 31-45). Such a modification would allow the method of Mohan to provide managing the input data for the database recovery independently of each other database by means of acquiring the log data by every database, in addition, the log data of the database to be recovered can be retrieved for a shorter time, allowing reduction of the time for the database recovery (see Tada col. 18, lines 35-45), therefore, improving the transaction performance and the accuracy of the parallel transaction recovery.

As per claim 32, Mohan discloses "recovering said plurality of dead transactions is performed by executing the particular number of recovery servers in parallel" (i.e., performing in parallel processing the new transactions; see page 722, col. 1, first paragraph, lines 7-9 and page 720, col. 1, second paragraph, lines 8-12).

As per claim 33, Mohan discloses "identifying the plurality of dead transactions includes the step of maintaining a working list" (i.e., updating and a log record is written; see page 719, col. 1, third paragraph, lines 1-5).

Mohan fails to explicitly disclose determining statistical data (statistic information) includes the step of determining statistical data (statistic information) based on the list dead transactions. However, Tada discloses a similar method, recovery operation on the database (see Tada col. 13, lines 16-22) determining statistical data (statistic information) includes the step of determining statistical data (statistic information) based on the list dead transactions (see Tada col. 13, line 65 to col. 14, line 2), further, Tada discloses the amount of occupation of the hlf buffer on nonvolatile memory. It would have been obvious to a person of ordinary skill in the art at the time the invention was made to modify the method of Mohan by determining statistical data (statistic information) includes the step of determining statistical data (statistic information) based on the list dead transactions as disclosed by Tada (see Tada col. 13, lines 31-45). Such a modification would allow the method of Mohan to provide managing the input data for the database recovery independently of each other database by means of acquiring the log data by every database, in addition, the log data of the database to be recovered can be retrieved for a shorter time,

allowing reduction of the time for the database recovery (see Tada col. 18, lines 35-45), therefore, improving the transaction performance and the accuracy of the parallel transaction recovery.

As per claim 34, in addition to claim 3, Mohan further discloses "locating a rollback segment, wherein the rollback segment contains a transaction table that contains entries associated with dead transactions" (i.e., missing updates transactions are redone during the redo pass; see page 720, col. 2, second paragraph, lines 1-5 and Fig. 3);

"scanning the transaction table to identify the dead transactions; and storing the identity of the dead transactions in the working list" (in light the specification at paragraph [0117], the purposed of scanning the transaction table is for identifying entries being associating with dead transactions. The method for using a log sequence number on every database page to track the page's state, in which every page update and log record is written is disclosed by Mohan page 719, third paragraph).

As per claim 35, in addition to claim 1, Mohan further discloses "maintaining a block count, wherein the block count identifies the number of undo blocks that are associated with a particular transaction; and the step of wherein the total number of undo blocks is based on the block count associated with the dead transactions identified in the working list" (In light the specification at paragraph [0052], the purposed of counting blocks is for associating a particular entry (data) in the transaction table. The method for tracking database page and placing in the field updated page, in which allowing precisely tracking restarting is disclosed by Mohan page 719, col. 2, paragraph 3).

As per claim 36, Mohan discloses "determining that the particular number of recovery servers should be used includes the step of determining that the particular number of recovery servers should be used based on a max_parallelism threshold value, wherein the max_parallelism threshold value provides an upper limit for the number of recovery servers to be used" (i.e., continually taking the maximum of the LSN (maximum parallel threshold) of the next log record to be processed for each of the yet to be

completely undone loser transaction until no loser transaction remains to be undone; see page 720, col. 1, second paragraph, lines 1-7 and Fig. 1).

As per claim 37, Mohan discloses "identifying a rollback segment that was previously owned by the crashed instance at the time of its crash" (i.e., updates rolling back; see page 720, col. 1, second paragraph, lines 1-2); and

"the crashed instance reacquiring ownership of the rollback segment after the crashed instance is restarted" (i.e., updating logged on stable storage but whose effects on the data base pages did not get reflected on disk before the system failure; see page 720, col. 1, first paragraph, lines 1-9).

As per claim 38, Mohan further discloses "identifying a rollback segment that is unowned, wherein the unowned rollback segment is not currently associated with any instance of the database" (i.e., identifying the transactions that are active, and addressing most recently written log; see page 719, col. 2, last paragraph, lines 1-6); and

"associating the unowned rollback segment with the crashed instance" (i.e., repeating failures during restarting recovery, which undo the same non record multiple times; see page 719, col. 2, last paragraph, lines 10-19), "wherein associating the unowned rollback segment with the crashed node causes the rollback segment to be owned by the crashed instance" (i.e., updates rolling back; see page 720, col. 1, second paragraph, lines 1-2).

As per claim 39, Mohan further discloses "recovering the plurality of dead transactions comprises the steps of maintaining a working list, wherein the working list identifies a list of dead transactions for which recovery will be attempted" (i.e., updating and a log record is written; see page 719, col. 1, third paragraph, lines 1-5);

"selecting a dead transaction from the working list" (i.e., selecting missing updates, loser transactions; see page 720, col. 1, paragraph 1, lines 4-11);

"acquiring a rollback segment lock on a rollback segment, wherein the rollback segment is associated with a transaction table that contains an entry that corresponds to the dead transaction" (i.e., performing update of the log record when locking is in effect; see page 722, col. 1, last paragraph to col. 2, up to line 9 and Fig. 6);

"acquiring a transaction lock on a chain of undo, wherein the chain of undo contain change information associated with the dead transaction" (i.e., performing update of the log record when locking is in effect; see page 722, col. 1, last paragraph to col. 2, up to line 9 and Fig. 6);

"determining whether the dead transaction still needs to be recovered; and if the dead transaction still needs to be recovered, assigning the dead transaction to a recovery server" (i.e., tracking of the progress of rollback; see page 719, col. 2, first paragraph, lines 10-17).

As per claim 40, in addition to claim 1, Mohan further discloses "assigning two or more dead transactions to a recovery server" (i.e., tracking of the progress of rollback; see page 719, col. 2, first paragraph, lines 10-17);

"associating a time slice value with the recovery server, wherein the time slice value is used by the recovery server to promote fairness during recovery of the two or more dead transactions; and recovering the two or more dead transactions using the time slice value" (In light the specification at paragraph [0151], the purposed of associating a time slice value with the recovery server is for recovering two or more dead transactions. The method for processing transaction continually until no loser (dead) transaction remains to be undone is disclosed by Mohan see page 720, col. 1, second paragraph, lines 1-10).

Examiner's Remark(s)

MPEP 2111: During patent examination, the pending claims must be "given the broadest reasonable interpretation consistent with the specification" Applicant always has the opportunity to amend the claims during prosecution and broad interpretation by the examiner reduces the possibility that the claim, once issued, will be interpreted more broadly than is justified. In re Prater, 162 USPQ 541,550-51 (CCPA 1969). The court found that applicant was advocating ... the impermissible importation of subject matter from the specification into the claim. See also In re Morris, 127 F.3d 1048, 1054-55, 44 USPQ2d 1023, 1027-28 (Fed. Cir. 1997) (The court held that the PTO is not required, in the course of prosecution, to interpret claims in applications in the same manner as a court would interpret claims in an infringement suit. Rather, the "PTO applies to verbiage of the proposed claims the broadest reasonable meaning of the words in their ordinary usage as they would be understood by one of ordinary skill in the art, taking into account whatever enlightenment by way of definition or otherwise that may be afforded by the written description contained in application's specification.").

The broadest reasonable interpretation of the claims must also be consistent with the interpretation that those skilled in the art would reach. In re Cortright, 165 F.3d 1353, 1359, 49 USPQ2d 1464, 1468 (Fed. Cir. 1999).

Prior Art

The prior art made of record and not relied upon is considered pertinent to applicant's disclosure.

Bamford et al., U.S. Patent. Number 5,832,516 relates to a method and apparatus for caching data in recoverable objects.

CONTACT INFORMATION

2. Any inquiry concerning this communication or earlier communications from the examiner should be directed to JEAN B. FLEURANTIN whose telephone number is 571 – 272-4035. The examiner can normally be reached on 7:05 to 4:35.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, JOHN E BREENE can be reached on 571 – 272-4107. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free).



Jean Bolte Fleurantin

Patent Examiner

Technology Center 2100

September 26, 2006